Agent-causal libertarianism, statistical neural laws and wild coincidences

Jason D. Runyan

forthcoming in *Synthese*

Abstract

Agent-causal libertarians maintain we are irreducible agents who, by acting, settle matters that aren't already settled. This implies that the neural matters underlying the exercise of our agency don't conform to deterministic laws, but it does not appear to exclude the possibility that they conform to statistical laws. However, Pereboom (1995; 2001; 2013) has argued that, if these neural matters conform to either statistical or deterministic physical laws, the complete conformity of an irreducible agent's settling of matters with what should be expected given the applicable laws would involve coincidences too wild to be credible. Here, I show that Pereboom's argument depends on the assumption that, at times, the antecedent probability certain behavior will occur applies in each of a number of occasions, and is incapable of changing as a result of what one does from one occasion to the next. There is, however, no evidence this assumption is true. The upshot is the wild coincidence objection is an empirical objection lacking empirical support. Thus, it isn't a compelling argument against agent-causal libertarianism.

1. Introduction

Agent-causal libertarians think we are irreducible agents who, by acting, cause certain changes, including bodily changes, and thereby settle matters that aren't already settled (e.g., Steward 2012). This implies that the neural matters underlying the exercise of our agency don't conform to deterministic laws, but it doesn't appear to exclude the possibility that they conform to statistical laws (cf. Clarke 1993, 2003; O'Connor 2000, 2008). However, Derk Pereboom (1995; 2001; 2013) has argued that, if these neural matters conform to either statistical or deterministic physical laws, the complete conformity of an irreducible agent's settling of matters with what should be expected given the applicable laws would involve coincidences too wild to be credible (Pereboom 2001, 85). Pereboom maintains that this objection, which I call 'the wild coincidence objection', is the most 'significant empirical objection' to agent-causal libertarianism (2001, 69). The reason is—on his view—our best scientific theories support the idea that neural matters conform to either deterministic or statistical laws.

Here, I show that the problem Pereboom elucidates depends on the assumption that, at times, the antecedent probability certain behavior will occur applies in each of a number of occasions, and is incapable of changing as a result of what one does from one occasion to the next. There is, however, no evidence this assumption is true. Rather, current findings from neuroscience would suggest that the antecedent probability that certain behavior will occur fluctuates from one occasion to the next, and is subject to change depending on what one does. The upshot is the wild coincidence objection is an empirical objection lacking empirical support. Thus, given the current state of neuroscience, it isn't a compelling argument against agent-causal libertarianism.

In what follows, I first outline initial reasons for thinking agent-causal libertarianism is compatible with the theory that the neural matters underlying the exercise of our agency conform to statistical laws (Section 2). I then present the wild coincidence objection (Section 3) before developing the case that it is an empirical objection without empirical support (Sections 4-5).

2. Agent-causal libertarianism and stochastic models

Agent-causal libertarians maintain we are agents *qua* substances—or persistent entities—who, by acting, cause change, and that our doing so can't be reduced to causings by mental or neural events and/or states (e.g., Steward 2012). They also maintain that, when we cause certain changes, we settle certain previously unsettled matters (e.g., Steward 2012, 34; Runyan 2016). Restated, we fix some of what obtains. For example, on a given occasion, an agent raises her arm, thereby causing her arm to rise and settling whether her arm rises on the occasion in question—a matter that hadn't yet been settled. In other words, the agent, by acting, fixes that her arm rises. In this sense, agent-causal libertarians maintain that, as agents, we are *irreducible settlers*.

The idea that we are irreducible settlers is incompatible with the idea that *all* neural matters conform to deterministic laws. *Deterministic laws* are laws that, along with the past, fix present and future matters. Thus, if all neural matters conform to deterministic laws, our nervous system is a *deterministic system*; or, that is, a system whose present and future behavior at any time *t* is fixed by what has already obtained and the relevant laws. And if our nervous system is a deterministic system, all neural matters leading to bodily motion when we act are settled, or fixed, *before* we do anything. We never settle a previously unsettled matter. We never fix anything that obtains by acting. In this case, we aren't irreducible settlers. Agent-causal libertarianism is false.

It has, however, been maintained that agent-causal libertarianism is consistent with the idea that the neural matters underlying the exercise of our agency conform to statistical, rather than deterministic, laws (cf. O'Connor 2000; Clarke 2003; Steward 2012, 244-45). Statistical laws are laws that specify the *probability* that certain matters will obtain on a given occasion; or on any one of a number of occasions. These laws apply to *indeterministic systems*; or, that is, systems whose behavior can only be probabilistically predicted, even when all the relevant variables are completely accounted for.

Predictive statistical models for the behavior of indeterministic systems take this general form:

(S1) The system's behavior ~ all the relevant variables + ε

In this schematic, predictors (variables used to predict) are on the right, and the criterion (what is predicted) is on the left (cf. Harrell 2001; Luke 2004). The ε—an error term—represents the extent to which the system's behavior deviates from the best predictions that can be generated from a *complete* and *accurate* account of all the relevant variables. The more a system's behavior deviates from what is predicted, the greater the value of ε. Accurate and complete models for the relationship between the behavior of indeterministic systems and all the relevant variables will always contain an error term. Accurate and complete models for the relationship between the behavior of deterministic systems and all the relevant variables will not. Given an accurate and complete model, the behavior of deterministic systems can be predicted without error.

If our functioning nervous systems are indeterministic systems, and statistical laws apply to at least some of their behavior, then schematic (S1), above, is an accurate and complete (though general) model of the relationship between a functioning nervous system's behavior and *all* the *neural* variables at play. By 'behavior', I mean patterns of neural activity and rest throughout a nervous system over any given period of time. Given we are agents, this would, at times, include the neural activity and rest underlying the exercise of our agency. By 'neural variables', for my purposes here I mean all variables relevant to a nervous system's behavior *excluding* anything a person *qua* irreducible agent does.[1]

Given schematic S1 is an accurate and complete model, even if we can precisely account for *all* neural variables at play, our nervous system's behavior will still deviate

[1] Likewise, by 'neural factors' I mean all factors relevant to a nervous system's behavior excluding anything the person *qua* irreducible agent does.

from the best predictions that can be generated using these variables alone. The most complete statistical model for the relationship will contain an error term (ε). Thus, it will be a *stochastic,* or probabilistic, model.

A potential explanation for the error term in a stochastic model for our nervous system's behavior would be that the model doesn't account for *all* factors influencing this behavior. There is at least one missing variable. And, if *all neural* variables have been accounted for in the model, the possibility arises that one variable not accounted for is the influence an agent *qua* irreducible settler has on her own nervous system. Thus, it seems at least *prima facie* that stochastic models are compatible with agent-causal libertarianism. For example, if an irreducible settler, by raising her arm, causes her arm to rise, and thereby settles whether it rises on a certain occasion, then she, likewise, settles whether certain changes occur throughout her nervous system leading to her arm rising. Thus, we should expect an error term in models for the relationship between the behavior of a person's nervous system and all neural variables at play. We should expect an error term in models that don't take into account what the agent *qua* irreducible settler does. In this case, though there are neural factors that increase, or decrease, the probability an irreducible settler will behave one way or another, how she behaves on a given occasion is settled, or fixed, by her when she acts or refrains; and not beforehand. Thus, we might say these neural factors are *inclining factors*. They incline the agent to behave in one way or another.

3. The wild coincidence objection

So far we've seen agent-causal libertarianism is incompatible with the idea that all neural matters conform to deterministic laws. We've also seen initial reasons for thinking agent-causal libertarianism is compatible with the idea that the neural matters underlying the exercise of our agency conform to statistical laws. However, this compatibility has come under dispute as a result of empirical objections developed by Pereboom.

While ‘agnostic about the truth of determinism’, Pereboom is a *hard incompatibilist*. That is, he maintains that, regardless of whether determinism is true, ‘we are not free in the sense required for moral responsibility’ (2013, 85). Pereboom argues that we can only have this sort of freedom if agent-causal libertarianism is true (e.g., 1995; 2001, 89-126; 2013, 85-111). At the same time, he also argues that—though ‘coherent’ and involving ‘no detected logical inconsistency’—there are ‘empirical considerations that tell against this position’ (2001, 68). On his view, agent-causal libertarianism faces ‘significant empirical objections’ (2001, 69; 2013, 111-14); the most significant of which is its inability to ‘accommodate our best natural science theories’ (2001, 69).

As Pereboom explains, if agent-causalism is true—and we are irreducible agents—there are agent-causal powers which aren’t powers had by our nervous system or any other part of our bodies (70-79). At the same time, according to Pereboom, given our best scientific theories, our nervous system’s behavior is ‘wholly governed by’ physical laws, whether deterministic or statistical. Thus, it would seem our exercise of the power to settle matters must either be constrained by these physical laws, or require our exercise of ‘the power to produce deviations from’ these laws; i.e.—‘deviations from what these laws would predict and from what we would expect given these laws’ (73, 79).

Two replies seem open to the agent-causal libertarian. The first is that, in order to exercise the power to settle matters, an irreducible agent needn’t produce deviations from what would be predicted given the physical laws (79). Pereboom calls this reply ‘reconciliation’. The second reply, which Pereboom calls ‘overriding’, is that an irreducible agent actually does ‘produce such deviations’ when she exercises the power to settle matters. In this case, our best scientific theories are wrong: our nervous system’s behavior isn’t ‘wholly governed by’ either statistical or deterministic physical laws.

While Pereboom thinks neither of these replies is 'credible', in the end Pereboom thinks 'overriding' provides the best bet for agent-causal libertarians (86). He arrives at this conclusion by presenting what he believes to be a decisive objection to the credibility of 'reconciliation'. The objection is that reconciliation would involve 'coincidences too wild to be believed' (2013, 112).

As Pereboom observes, given agent-causal libertarianism:

> That the agent-cause is a causal factor distinct from the causal factors that incline her is underscored by her capacity to act in opposition to them. For instance, even if her reasons incline her very strongly toward performing an action at each opportunity for performing it, she can choose to refrain every time. But what mechanism could then explain the agent-cause's conforming, *in the long run*, to the same frequency of choices that would be extremely likely to obtain on the basis of the inclining factors alone? On the agent-causal view, if the agent-cause is truly free, there is no mechanism that could provide this explanation. We would therefore have a match of frequencies without an explanation—a wild coincidence. (2001, 84-5; my emphasis)

Pereboom's objection is meant to apply irrespective of whether the inclining factors are psychosocial and/or neurophysiological factors (84; cf. 2013, 112). The point is, due to inclining factors, there will be a certain probability an agent will behave in a certain way given the opportunity. Thus, we should expect this behavior to occur with a certain frequency across multiple occasions (cf. 2013, 112). In this case, certain statistical laws, rooted in inclining factors, apply on each of these occasions. And they apply *independent* of how the irreducible agent exercises her ability to settle matters. This independence is a result of the agent being a distinct cause, as attested to by her ability to act in opposition to inclining factors (2013, 113). Further, because of this independence, Pereboom thinks we should 'expect evidence of the effect of the additional cause, the agent-cause, to show up in the long run in the actual frequencies' (2001, 85; cf. 2013, 113). If, however, an irreducible agent's settling of matters conforms, 'in the long run', to what should be expected given the applicable statistical laws, there can be no explanation for this conformity. Over the long haul, the 'dovetail' of what should be expected given these laws and the pattern of an irreducible agent's settling of matters would involve 'coincidences so wild as to make it incredible' (2001, 85).

To illustrate the problem at hand, Pereboom (2013) provides an example:

> Consider the class of possible actions, each of which has a physical component whose antecedent probability of occurring is approximately 0.32. It would not violate the statistical laws in the sense of being logically incompatible with them if, for a large number of instances, the physical components in this class were not actually realized close to 32% of the time. Rather, the force of the statistical law is that for a large number of instances, it is correct to *expect* physical components in this class to be realized close to 32% of the time. Are free choices on the agent-causal libertarian model compatible with what the statistical law would lead us to expect about them? If they were, then for a large enough number of instances, the possible actions in our class would almost certainly be freely chosen nearly 32% of the time. But if the occurrence of these physical components were settled by the choices of agent-causes, then their actually being chosen close to 32% of the time would also amount to a wild coincidence. The proposal that agent-caused free choices do not diverge from what the statistical laws predict for the physical components of our actions would be so sharply opposed to what we would expect as to make it incredible. (112)

Here, Pereboom has us consider 'the class of possible actions, each of which has a physical component' that has a .32 probability of occurring *antecedent* to—and independent of—an irreducible agent settling whether she performs one of these actions. According to Pereboom, *if (or when)* this probability applies in each of 'a large enough number of instances', the applicable statistical laws are such that we should expect actions in this class to occur close to 32% of the time. The problem is, if—at the same time—the agent is truly exercising her ability to settle what to do across these instances, that she performs these actions close to 32% of the time must be a wild coincidence. There can be no explanation for it. Thus, given such cases, that expectations would always be met by an irreducible agent's settling of matters would seem to defy the odds to the point of being incredible.[2]

In reply to the wild coincidence objection, agent-causal libertarians might suggest, as Roderick Chisholm (1964/2002) has, that when we are in position to exercise our agency the underlying neural matters no longer conform to the statistical laws they otherwise do.[3] Or the agent-causal libertarian might even suggest, as Timothy O'Connor (2008) has, that underlying neural matters conform to some other kind of law—which is neither deterministic nor statistical—specific to the situations within which an agent *qua* irreducible settler exercises her agency. However, these replies seem *ad hoc* and lack evidence. Further, and more important for our purposes here, such replies amount to a concession that agent-causal libertarianism is implausible given the underlying neural matters conform to statistical laws.

[2] It should be noted that, while Pereboom is talking about a class of possible actions, he is more precisely talking about what would be the case if (or when) a .32 probability 'a physical component' of each of these actions will occur *actually applies* in each of 'a large enough number of instances' (2013, 112). This is made clear by his focus on what we should expect, 'in the long run' (2001, 85), if (or when) this is, in fact, the case; namely—that the actions in this class are chosen 'close to 32% of the time' across these instances. The reason for this focus is that—as Pereboom realizes and as will come into play in Section 4—there must be at least one action with a physical component that has a certain antecedent probability of occurring in each of a large enough set of instances for the 'dovetail' of expectations and an irreducible agent's settling of matters to 'amount to a wild coincidence' (2001, 85). If not, agent-causal libertarians can offer a plausible explanation for why expectations are always met. To illustrate, if, for example, a .32 probability an agent will perform any one of a number of actions applies to *only one* instance, the agent can perform any of these actions or none of them, and her behavior will conform to expectations either way. Thus, in such cases, expectations are always met regardless of what the agent does. I thank an anonymous reviewer for helping me see the importance of clarifying this here.

[3] M. Vargas (2013) has recently suggested agent-causal libertarians are likely to take this position (69).

4. Probabilities, inclining factors and change

Recently, Jordan Baker (2016) has argued that Pereboom's wild coincidence objection doesn't present a problem for agent-causal libertarians as long as they adhere to what he calls 'the integrated approach'. According to this approach, as a matter of natural law, whenever certain inclining factors produce certain effects, so do certain irreducible agent settlers—as 'co-producers' (cf. Clarke 2003, 145). Consequently, the antecedent probability physical components of certain actions will occur is also the probability an irreducible agent will settle matters in a certain way. So these probabilities don't apply independently. Thus, we should expect an irreducible agent's settling of matters to meet expectations given the antecedent probabilities.

This is all I will say about Baker's reply to the wild coincidence objection. The reason is—as I will argue here—there is a more fundamental problem facing the wild coincidence objection. Regardless of whether the integrated approach succeeds, the 'dovetail' of what should be expected given the applicable statistical laws and an irreducible agent's settling of matters needn't ever be a wild coincidence. To think otherwise is to make a problematic assumption about the *antecedent* probability an indeterministic system will behave a certain way on each of a number of occasions; namely—that this probability isn't subject to change depending on the system's behavior across these occasions. And this holds implications for the success of the wild coincidence objection as an *empirical* objection to agent-causal libertarianism. Let me start with an illustration.

If our nervous system's behavior conforms to statistical laws, it might be established in advance that, if an individual—let's call her Ellie—has the opportunity to raise her hand to answer a question in class, there is a 32% chance neural activities leading to her hand rising will occur.[4] Let's grant that this probability applies

[4] For illustrative purposes—to keep things simple—I will focus on a *particular* possible action (i.e., hand-raising) with a physical component that has a certain antecedent probability of occurring in each of a number of instances. This is a slight variation from Pereboom's example—discussed in Section 3—involving a *class* of such actions. However, in this context, the main difference between considering a specific action and a class of actions is the number of physical components involved, and the overall

independent of the probability she will exercise her agency by raising her hand. (So, here, we are assuming the integrationist approach isn't correct.) Given this, if this .32 probability applies on each of 100 occasions, then, in order for Ellie's behavioral patterns to conform to expectations given the applicable statistical laws, Ellie must raise her hand close to 32 times. The problem Pereboom has raised is that, if, at the same time, this hand-raising is an exercise of Ellie's ability to settle whether she raises her hand, it would seem that this conformity would amount to a wild coincidence. There would be no explanation for it, and it would defy the odds.

Here, it is important to note that statistical tests can help us more precisely establish the frequencies which meet expectations given certain probabilities. One of the most common is chi-square (e.g., Greenwood & Nulkin 1996). Performing a chi-square can help us more precisely establish the frequencies with which we should expect an individual to engage in a certain behavior given the antecedent probability the behavior will occur. So—returning to our example—if the antecedent probability Ellie's hand will rise on each of 100 occasions is .32, chi-square results indicate we should conclude Ellie's behavior meets expectations as long as she raises her hand anywhere between 25 (χ_1=2.25, p=.13) and 39 (χ_1=2.25, p=.13) times.[5]

Now, at this point, having a more precise handle on expectations doesn't seem to change much. Pereboom's objection seemingly still holds. It would still seem to involve a wild coincidence for Ellie—or any irreducible agent—to always, in the long run,

---

complexity. This shift also enables us to consider a particular example, and action, which helps keep things more concrete. For my purposes here, this shift is, otherwise, inconsequential. Whether talking about a physical component of a particular action or a class of actions—as we will see below—the crux of the matter is that there must actually be at least one possible action with a physical component that has a certain antecedent probability of occurring in each of 'a large enough number of instances' (Pereboom 2013, 112), and this probability *must* be incapable of changing as a result of what one does from one instance to the next, for the wild coincidence problem to work. Otherwise, the agent-causal libertarian can give a plausible explanation for why an irreducible agent's behavioral patterns *always* conform to expectations. In other words, this conformity needn't be a wild coincidence. This, for example, includes when this conformity involves all 'possible actions, *each of which* has a physical component whose antecedent probability of occurring is approximately 0.32' (Pereboom 2013, 112; emphasis added) as much as a particular possible action with a physical component that has a .32 probability of occurring. I thank an anonymous reviewer for bringing it to my attention that it would be helpful to point this out.

[5] Here, I will be conservative by using a *p*-value of .10 to determine whether it should be concluded that expectations aren't met. In comparison with the traditional value of .05, this will increase the likelihood of drawing the conclusion that expectations aren't met.

settle matters with a pattern that conforms to expectations given the antecedent probabilities. But, despite appearances, there is a credible explanation for this conformity.

To see this, for simplicity sake, let's stick with the example of Ellie, but focus on the first five occasions where she settles whether she raises her hand. At the outset, on each of these occasions—as is the case for the next 95—the *antecedent* probability Ellie's hand will rise is .32. Chi-square tells us that, across these five occasions, Ellie could raise her hand anywhere between zero (χ^2_1=2.35, p=.13) and three (χ^2_1=1.80, p=.18) times, and expectations would be met.[6] So, again, a variety of frequencies of behavior meet expectations. Additionally, recall that the probability one will behave in a certain way is rooted in inclining factors, which might include psychosocial and neurophysiological factors. Thus, inclining factors include such things as reasons, desires, urges and motivations as well as neurophysiological states and events (e.g., Pereboom 2001, 84). These inclining factors, by definition, incline an agent to behave in certain ways and not in others. The force of inclining factors—and any corresponding statistical laws—is that they make it more or less likely an agent will behave in a certain way. Consequently, certain behaviors may, for example, be more effortless, spontaneous or more motivated, while others more effortful, deliberate or less motivated.

Further, it is important to note that statistical laws and probabilities may, *at the outset*, apply to each occasion in a set, and yet be rooted in inclining factors that can wax or wane in force across these occasions.[7] To illustrate, these laws might be rooted in desires, urges, attitudes, motives, reasons or states, which can change across these occasions depending on what one does from one occasion to the next.[8] Equally, they

[6] Barnard's test—another test for whether frequency patterns meet expectations—also indicates Ellie could raise her hand anywhere between zero (Wald Stat.=1.52; Nuissance parameter=.01; p=.13) and three times (Wald Stat.=1.3, Nuissance parameter=.01, p=.15) and expectations would be met.

[7] In this case, even though, at the outset, one has a certain probability of behaving in a certain way in each of a set of occasions, the probability one will behave in this way might end up differing from one occasion to the next depending on what one does on each occasion.

[8] For example, there is psychological evidence that desires and will-power can wax and wane depending on what one does, or doesn't do, throughout the day (e.g., Hofmann 2011).

might be rooted in neurophysiological factors that may change depending on what one does. And, if the underlying inclining factors can wax and wane from one occasion to the next on the basis of what one does, so can the probability that certain behavior will occur. As a consequence, the waxing and waning of inclining factors might, at times, place parameters on how frequently an irreducible agent will exercise her agency in a particular way across a set of occasions. For instance, if eating is exercising one's ability to settle whether one eats, the waxing and waning of our appetite might put parameters on how frequently we eat across a set of occasions. This possibility not only seems quite credible, but to be expected.[9]

Thus—in the case of Ellie—certain inclining factors might wax, and others wane, across five occasions depending on what she does from one occasion to the next. And this might, at times, place parameters on how frequently she raises her hand. All this might be true in spite of the fact that, *at the outset*, given the applicable statistical laws, the probability Ellie's hand will rise on any one of these occasions is .32. Consequently, even if the antecedent probability Ellie's hand will rise on any one of five occasions is .32 *and expectations are met*, Ellie might—for example—raise her hand all three times across the first three occasions. If she does this, it might simply be the case that she will not raise it over the next two occasions due to the effects she has on inclining factors (e.g., her motivation).[10] Alternatively, Ellie might raise her hand on three out of the first four occasions. If she does this, it might be that she will not raise her hand on the next occasion (again, perhaps, due to a lack of motivation as a result of her previous actions). Or Ellie might raise her hand on the third and fifth occasions; and, as a result, the relevant inclining factors, and the probability Ellie will raise her hand, may *remain unchanged*.

In short, over the five occasions in question—though she might exercise her agency such that the probability her hand will rise remains the same—Ellie might

[9] For example, Steward (2012), an agent-causal libertarian, seems to suggest as much when she indicates that—on her view—agents sometimes settle matters in advance (cf. 39-42).

[10] Of course, this doesn't rule out the possibility that something might intervene and thereby affect inclining factors in a way that increases the probability she will raise her hand on these occasions. This might include other things she may do (e.g., engage in certain thought processes).

exercise her agency in various ways that influence the probability her hand will rise on the next occasion. Equally, *the pattern* with which she exercises her agency across this limited set of occasions might have *longer-term influences* going forward. For example, if Ellie doesn't raise her hand at all across the five occasions—which would be in keeping with expectations—inclining factors might wane across later occasions, and the probability she will raise her hand may decrease down the road. Conversely, if she raises her hand on three of the five occasions, inclining factors may increase going forward, as may the subsequent probability she will raise her hand. The point is the probability Ellie's hand will rise on *subsequent* occasions might be *subject to change* depending on what she does across the five occasions in question. And the same might be the case for the five occasions after that, and the five after that, etc. Thus, while the antecedent probability Ellie's hand will rise on any one of 100 occasions might be .32, this probability might be subject to change depending on what she does across these occasions. The probability her hand will rise may increase along with the frequency with which she raises it; or vice versa.

So we've seen that the *antecedent* probability Ellie's hand will rise on any one of a rather large number of occasions might be .32; and yet this probability might still be subject to change depending on what she does across these occasions—both in the short- and long-term. It is, of course, true that Elle wouldn't settle:

(i) that the antecedent probability her hand will rise is subject to these changes.

She also wouldn't settle:

(ii) that, at times, the waxing and waning of inclining factors places parameters on how frequently she will raise her hand across a set of occasions.[11]

[11] Though she might influence whether they wax or wane.

Nevertheless, (i) and (ii) might very well explain why Ellie's behavior *always* conforms with expectations given the relevant probabilities, *even while she settles whether she raises her hand on each occasion*.[12] As illustrated above, it is quite credible that even while Ellie exercises the ability to settle whether she raises her hand, (i) and (ii) make it so expectations are always met. And all this is true even in the event these probabilities *remain unchanged* across a number of occasions, which—given the force of inclining factors—should be expected with some frequency. In such cases, the explanation would simply be that Ellie exercised her agency with a pattern that conformed to the relevant probabilities, and thus in a way that didn't alter them.

As a brief aside, I should clarify that—under certain conditions—there needn't be parameters on how frequently an irreducible agent will exercise her agency in a particular way for there to be an explanation for why expectations are met no matter what she does. To illustrate, if there were a .32 probability Ellie would raise her hand on each of five occasions, and she raised it on the first occasion, this might make the probability jump to .65 on the second occasion due to changes in inclining factors. She might, again, raise her hand on this occasion making the probability jump further—this time to .7. She may, then, raise her hand on the third and fourth occasions, which might make the probability jump yet again to .75 and .87, respectively. The *mean* probability of her raising her hand on these occasions would end up being .658; and expectations would be met even if she raised her hand on each occasion (χ^2_1=2.60, p=.11). This would be the case even though the antecedent probability she would raise her hand on each of these five occasions had been .32.[13]

[12] On some occasions, that she doesn't raise her hand might be settled by her across previous occasions, where, by acting, she makes it so inclining factors wane to the point she has no motivation, and will not act on a subsequent occasion, barring some change.

[13] Alternatively, imagine there is a .87 probability Ellie will raise her hand on each of five occasions (cf. Pereboom 2001, 85). She may, nevertheless, not raise it on the first occasion, which might make the probability she will raise it on the second occasion drop to .25. Again, she may not raise her hand on this occasion, making the probability she will raise it on the third occasion drop to .2. She may, then, not raise her hand on the third and fourth occasions, which might make the probability drop further to .17 and .15, respectively. Given this, the mean probability of her raising her hand on these occasions ends up being .328; and expectations are met even if she never raises her hand (χ^2_1=2.44, p=.12).

At this point, it might be helpful to summarize why it is that—given (i) and (ii), above—there may very well be an explanation for why Ellie's behavior *always* conforms with expectations. The explanation would—in general terms—have two parts. First, inclining factors may wax and wane depending on what Ellie does. This, at times, sets parameters on how frequently she will exercise her agency in certain ways across a set of occasions. Second, Ellie either exercises her agency in a way that shifts expectations going forward (when these expectations change), or in a way that conforms with expectations and doesn't change them (when they don't change). She settles which is the case, but either way—and *regardless* of what she does—expectations are met. Expectations are always met since the patterns with which Ellie exercises her agency are, to some degree, restricted, and yet she has an influence on expectations going forward. It is this combination of restrictions and influence which might very well explain why, *in every case*, her behavior conforms with expectation. As a result, this conformity needn't ever happen by wild coincidence.[14]

The upshot is—even if we put aside the integrated approach—there may very well be an explanation as to why an irreducible agent's settling of matters always conforms with what should be expected given the applicable statistical laws. As we've seen here, that there would be such an explanation becomes plausible *if* the relevant probabilities are subject to change depending on what the agent does from one occasion to the next. In this case, the conformity in question doesn't necessarily involve wild coincidences. Consequently, the wild coincidence objection depends on it being true that, at times, the antecedent probability certain behavior will occur applies in each of a number of occasions, and is *incapable* of changing as a result of what one does from one occasion to the next. The problem is Pereboom offers no reason to think this is ever the case. And, given the wild coincidence objection assumes our nervous system might be an *indeterministic* system (i.e., a system to which statistical laws apply), we can't simply assume that, at times, the antecedent probability that certain behavior will occur is *incapable* of changing in this manner. Further, as we will see next, the available

[14] I thank an anonymous reviewer for alerting me to the value of providing this summary.

neuroscientific evidence doesn't support such an assumption. This is problematic for the wild coincidence objection since it is an empirical objection to agent-causal libertarianism.

5. The wild coincidence objection: an empirical objection without empirical support

So we've seen that the wild coincidence objection depends on the assumption that, at times, the antecedent probability certain behavior will occur applies in each of a number of occasions, and is incapable of changing as a result of what one does from one occasion to the next. As we will now see, there isn't any evidence this is the case. Rather, over the past several years, there has been accumulating evidence that the antecedent probability motor-related neural activities—and, consequently, certain behavior—will occur on a given occasion is subject to change depending on what one does on previous occasions. In fact, multiple lines of neuroscience research are beginning to capture, and quantify, the extent of these activity-dependent changes.[15] Of particular relevance here, it has been observed that making a spontaneous choice on one occasion increases predictive neuronal activity patterns for making the same choice on the next occasion (e.g., Bode et al. 2011, 2013, 2014; Lavazza 2016). Additionally, in a task where one is repeatedly forced to make a choice between two options (i.e., spontaneously move one's right or left hand), it has been observed that how one chooses on a given trial influences the probability one will make the same choice on the next trial (Soon et al. 2014).

Further, it is well-known that activity-dependent neural changes reflect forms of associative (or Hebbian) learning, whereby interconnected neurons that are active

[15] For e.g., see: Wolters et al., 2003; Huber et al. 2004; Floyer-Lea et al. 2006; Roy et al. 2007; Albert et al. 2009; Feldman 2009; 2012; Wang 2010; Mrachacz-Kersting 2012; Orban de Xivry et al. 2013; Koch et al. 2013; Hammerbeck et al. 2014; Barker et al. 2014; McNickle & Carson 2015; Chao et al. 2015. Also, it has, for instance, been observed that, as a behavior is repeated, certain motor-related neural activities become more probable under certain conditions (e.g., Costa 2007; Wickens et al. 2007; Graybiel 2008; Verstynen & Sabes 2011; Hikosaka et al. 2013; Kim et al. 2015; Anderson 2016).

together are subsequently more likely to be active together under certain conditions, and vice versa (e.g., Hebb 1949; Klaes et al. 2012). These changes can happen rapidly (within milliseconds), and involve intracellular biochemical changes which influence synaptic transmission among other things (e.g., Friedman et al. 2015; Cichon and Gan 2015; Svensson et al. 2016). They can also involve rapid morphological changes (within minutes to hours) to neurons and their synaptic connections (e.g., Kandel 2001; Dash et al. 2004; Fu et al. 2012; Bailey et al. 2015; Chen et al. 2015; Hayashi-Takagi et al. 2015). The ability of nervous systems, and their states, to be changed in this way is called *plasticity*; and plasticity generalizes across functioning nervous systems (e.g., Fu et al. 2002; Jacob et al. 2007; Cassenaer & Laurent 2007; Yin et al. 2009; Nazzaro et al. 2012; Smith & Graybiel 2013; Arce-McShane et al. 2014). Nervous systems, and their subsystems, are 'plastic' in the sense that their proper functioning throughout daily life involves changes in the probability subsequent activities will occur (e.g., Dan & Poo 2006; McMahon & Leopold 2012; Veniero et al. 2013). For instance, plasticity-related changes have been observed to occur from morning to evening, within a single day (Shannon et al. 2013). In particular, changes in motor cortical (M1) excitability have been observed to occur as a result of performing, or even watching or imagining, simple movement tasks (e.g., McCombe Waller et al. 2008; Kumru et al. 2008, 2016).

To add to the above, there is growing evidence that, in conscious animals and humans at rest, coordinated cortical activity, called 'avalanches', organize from spontaneous fluctuations, and shape the probability of future cortical activity patterns (cf. Fox and Raichle 2007; Luczak et al. 2009; Palva et al. 2013; Bellay et al. 2015). In keeping with this, there is evidence that preparatory neural activity patterns, which increase the likelihood of a spontaneous choice to act, are influenced by the ebb and flow of stochastic fluctuations (cf. Schurger et al. 2012, 2015, 2016). Thus, the available evidence suggests that the probability motor-related neural activities will unfold—and that we will behave—in certain ways changes from one occasion to the next due to spontaneous neural fluctuations occurring while we are at rest.[16] And this would be in addition to any changes in probability that may result from what we have previously done (cf. Lavazza 2016).

In sum, the available neuroscientific evidence suggests that—when we are conscious, and potentially have the opportunity to exercise our agency—the probability that certain behavior will occur fluctuates from one occasion to the next; and that what we do has an influence on subsequent probabilities concerning how we will behave in the future. If this turns out to be the case, the probability one will do something in a given situation should be thought of as existing on a sliding scale (see Figure 1). It can slide, or change, as one goes through life from one occasion to the next. This is a problem for the wild coincidence objection. The reason is—as observed in Section 4—this objection depends on the assumption that, at times, the antecedent probability certain behavior will occur applies in each of a number of occasions, and is *incapable* of changing as a result of what one does from one occasion to the next. Thus, as an empirical objection to agent-causal libertarianism, the wild coincidence objection requires some evidence of this to be convincing.

[16] While nothing I've said here depends on it, it is tempting to speculate that these spontaneous fluctuations might increase the likelihood of variability in an agent's choices and, thus, in her experiences. Engaging in a variety of different behaviors, and having a variety of experiences, in the same or similar circumstances would increase an agent's exposure to novelty. This may have advantageous consequences on flexibility, increased awareness of possibilities, and learning. I thank an anonymous reviewer for prompting me to make this point.

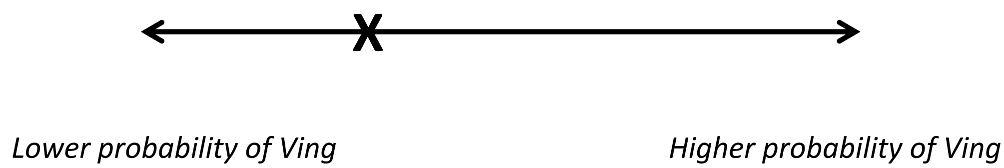

Figure 1. *A representation of a sliding scale for the probability one will behave in a certain way (V) in a given situation.* The line represents a range of possible probabilities one will behave in a certain way in a given situation. The probability one will behave in a certain way on a given occasion, represented by the **X**, may shift depending on what one does across previous occasions.

Before drawing this section to a close, I should mention that it is hard to see how there could be evidence to support the wild coincidence objection. In keeping with what we saw in Section 4, it is plausible that the probability one will behave in a certain way on any one of a number of occasions will remain close to being the same *unless* one's behavioral patterns *deviate* from the patterns expected *given this probability doesn't change*. And one is likely to often exercise one's agency with a pattern that doesn't deviate from these expectations—this is the force of the inclining factors from which these expectations stem. Consequently, if our nervous system is an indeterministic system,[17] it becomes hard to see how we might identify cases where the relevant probabilities are *incapable* of being changed over multiple occasions. It would seem difficult to rule out the possibility these probabilities *simply don't* change while being capable of it. In this case, since the wild coincidence objection depends on it being true that, at times, these probabilities are incapable of changing as a result of what one does from one occasion to the next, it is hard to see how there could be evidence in favor of the wild coincidence objection.

[17] And, again, the wild coincidence objection involves assuming this might be the case.

6. Conclusions

In conclusion, given what we know about nervous systems, we shouldn't find the wild coincidence objection convincing as an empirical objection to agent-causal libertarianism. If our nervous system's behavior conforms to statistical laws, we shouldn't assume that, at times, the antecedent probability that certain behavior will occur applies in each of a number of occasions, or is incapable of changing as a result of what one does from one occasion to the next. Rather—in keeping with observations made in Section 5—the available evidence suggests that, across occasions where one might exercise one's agency, the probability that certain behavior will occur might very well change depending on what one does, as well as due to spontaneous neural fluctuations. Consequently—as illustrated in Section 4—there might be a credible, and even expected, explanation as to why an irreducible agent's settling of matters always conforms with what should be expected given the applicable statistical laws. This conformity needn't involve wild coincidences. Instead, it might plausibly be explained by the fact:

- the probability that certain behavior will occur on a given occasion is subject to change depending on what an irreducible agent does across previous occasions;

along with the fact that:

- at times, the waxing and waning of inclining factors places parameters on how frequently an irreducible agent exercises her agency in a particular way across a set of occasions.

While neither point may be an established fact, each is credible and there doesn't seem to be any evidence against either. If anything, these points are supported by our current neuroscientific findings.

The upshot is the wild coincidence objection is an empirical objection lacking empirical support. Thus, it isn't a compelling objection to agent-causal libertarianism. This observation is significant if, as Pereboom has argued: (i) we only have the kind of freedom required for moral responsibility if agent-causal libertarianism is true; (ii) agent-causal libertarianism is coherent; and (iii) the wild coincidence objection serves as the most significant empirical objection to agent-causal libertarianism. The reason is, if Pereboom is right on all three accounts, in as much as we are agnostic about the truth of determinism, we shouldn't rule out the idea we have the kind of freedom required for moral responsibility.

References

Anderson, B. A., Kuwabara, H., Wong, D. F., Gean, E. G., Rahmim, A., Brašić, J. R., George, N., Frolov, B., Courtney, S. M., Yantis, S. (2016). The Role of Dopamine in Value-Based Attentional Orienting. *Current Biology* 26: 550-5.

Albert N. B., Robertson E. M., Miall R. C. (2009). The resting human brain and motor learning. *Current Biology* 19: 1023–1027.

Alvarez, M., Hyman, J. (1998). Agents and their actions. *Philosophy* 73: 219-45.

Arce-McShane, F. I., Hatsopoulos, N. G., Lee, J. C., Ross, C. F., Sessle, B. J. (2014). Modulation dynamics in the orofacial sensorimotor cortex during motor skill acquisition. *Journal of Neuroscience* 34: 5985-97.

Baily, C. H., Kandel, E. R., Harris, K. M. (2015). Structural components of synaptic plasticity and memory consolidation. *Cold Spring Harbor Perspectives in Biology* 7. doi: 10.1101/cshperspect.a021758.

Baker, J. (2016). Rejecting Pereboom's empirical objection to agent-causation. *Synthese*. doi: 10.1007/s11229-016-1094-0.

Bellay T, Klaus A, Seshadri S, Plenz D. (2015). Irregular spiking of pyramidal neurons organizes as scale-invariant neuronal avalanches in the awake state. *Elife* 4: e07224.

Barker, J. M., Taylor, J. R., Chandler, L. J. (2014). A unifying model of the role of the infralimbic cortex in extinction and habits. *Learning and Memory* 21: 441-8.

Bode, S., He, A. H., Soon, C. S., Trampel, R., Turner, R., and Haynes, J. D. (2011). Tracking the unconscious generation of free decisions using ultra-high field fMRI. *PLoS One* 6: e21612. doi: 10.1371/journal.pone.0021612.

Bode, S., Bogler, C., and Haynes, J. D. (2013). Similar neural mechanisms for guesses and free decisions. *Neuroimage* 65: 456–465.

Bode, S., Murawski, C., Soon, C. S., Bode, P., Stahl, J., and Smith, P. L. (2014). Demystifying "free will": the role of contextual information and evidence accumulation for predictive brain activity. *Neuroscience & Biobehavioral Reviews* 47, 636–645.

Cassenaer, S., Laurent, G. (2007). Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* 448: 709–713

Chao C. C., Karabanov A. N., Paine R., Carolina de Campos A., Kukke S. N., Wu T., Wang H., Hallett M. (2015). Induction of motor associative plasticity in the posterior parietal cortex-primary motor network. *Cerebral Cortex* 25: 365-73.

Chen S.X., Kim A.N., Peters A.J., Komiyama T. (2015). Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. *Nature Neuroscience* 18: 1109-15.

Chichon, J., Gan, W. B. (2015). Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity. *Nature* 520: 180-5.

Chisholm, R. (1964/2002). 'Human freedom and the self', in *Free Will*, ed. R. Kane. Oxford: Blackwell.

Clarke, R. (1993). Toward a credible agent-causal account of free will. *Noûs* 27: 191-203.

Clarke, R. (2003). *Libertarian accounts of free will*. Oxford: Oxford University Press.

Costa, R. M. (2007). Plastic Corticostriatal Circuits for Action Learning: What's Dopamine Got to Do with It?. *Annals of the New York Academy of Science* 1104: 172-91.

Dan, Y., Poo, M. M. (2006). Spike-timing-dependent plasticity: from synapse to perception. *Physiology Review* 86: 1033–1048.

Feldman, D.E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annual Review of Neuroscience* 32: 33–55.

Floyer-Lea, A., Wylezinska, M., Kincses, T., Matthews, P. M. (2006). Rapid modulation of GABA concentration in human sensorimotor cortex during motor learning. *Journal of Neurophysiology* 95: 1639-44.

Fox M. D., Raichle M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience* 8:700-11.

Friedman A. K., Weiss, K. R., Cropper, E. C. (2015). Specificity of repetition priming: the role of chemical coding. *Journal of Neuroscience* 35: 6326-34.

Fu M, Yu X, Lu J, Zuo Y. (2012). Repetitive motor learning induces coordinated formation of clustered dendritic spines *in vivo*. *Nature* 483: 92-5.

Fu, Y.X., Djupsund, K., Gao, H., Hayden, B., Shen, K., Dan, Y. (2002). Temporal specificity in the cortical plasticity of visual space representation. *Science* 296: 1999–2003.

Greenwood, P.E., Nikulin, M.S. (1996). *A guide to chi-squared testing*. New York: Wiley.

Graybiel, A. M. (2008). Habits, rituals, and the evaluative brain. *Annual Review of Neuroscience* 31: 359-87.

Hammerbeck, U., Yousif, N., Greenwood, R., Rothwell, J. C., Diedrichsen, J. (2014). Movement speed is biased by prior experience. *Journal of Neurophysiology* 111: 128-34.

Harrell, Jr., F. E. (2001). *Regression Modeling Strategies: With applications to linear models, logistic regression, and survival analysis*. New York: Springer.

Hayashi-Takagi, A., Yagishita, S., Nakamura, M., Shirai, F., Wu, Y.I., Loshbaugh, A.L., Klaes, C., Schneegans, S., Schöner, G., Gail, A. (2012). Sensorimotor learning biases choice behavior: a learning neural field model for decision-making. *PLoS Computational Biology* 8: e1002774l

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley & Sons.

Hikosaka, O., Yamamoto, S., Yasuda, M., Kim, H. F., (2013). Why skill matters. *Trends in Cognitive Science* 17: 434-41.

Hofmann, W., Baumeister, R. F., Förster, G., Vohs, K. D. (2011). Everyday temptations: An experience sampling study of desire, conflict, and self-control. *Journal of Personality and Social Psychology* 102: 1318-35.

Huber R., Ghilardi M. F., Massimini M., Tononi G. (2004). Local sleep and learning. *Nature* 430: 78–81.

Jacob, V., Brasier, D. J., Erchova, I., Feldman, D., Shulz. D.E. (2007). Spike-timing-dependent synaptic depression in the in vivo barrel cortex of the rat. *Journal of Neuroscience* 27: 1271–1284

Kandel, E. (2001). The molecular biology of memory storage: A dialogue between genes and synapses. *Bioscience Reports* 21: 565-611.

Kane, R. (2005). *A contemporary introduction to free will*. New York: Oxford University Press.

Klaes C, Schneegans S, Schöner G, Gail A. (2012). Sensorimotor learning biases choice behavior: a learning neural field model for decision making. *PLoS Computational Biology* 8: e1002774

Kim, H. F., Ghazizadeh, A., Hikosaka, O. (2015). Dopamine Neurons Encoding Long-Term Memory of Object Value for Habitual Behavior. *Cell* 163: 1165-75.

Koch, G., Ponzo, V., Di Lorenzo, F., Caltagirone, C., Veniero, D. (2013). Hebbian and anti-Hebbian spike-timing-dependent plasticity of human cortico-cortical connections. *Journal of Neuroscience* 33: 9725-33.

Kuhlman, B., Hahn, K.M., Kasai, H. (2015). Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* 525: 333-8.

Kumru H, Soto O, Casanova J, Valls-Sole J. (2008). Motor cortex excitability changes during imagery of simple reaction time. *Experimental Brain Research* 189: 373-8.

Kumru H, Albu S, Pelayo R, Rothwell J, Opisso E, Leon D, Soler D, Tormos JM. (2016). Motor cortex plasticity during unilateral finger movement with mirror visual feedback. *Neural Plasticity* 2016: 6087896.

Lavazza, A. (2016). Free Will and Neuroscience: From Explaining Freedom Away to New Ways of Operationalizing and Measuring It. *Frontiers in Human Neuroscience* 10. Doi: 10.3389/fnhum.2016.00262.

Luczak A, Barthó P, Harris KD. (2004). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 2009 62: 413-25.

Luke, D. E. (2004). *Multilevel modeling*. London: Sage.

McCombe Waller S., Forrester L., Villagra F., Whitall J. (2008). Intracortical inhibition and facilitation with unilateral dominant, unilateral nondominant and bilateral movement tasks in left- and right-handed adults. *Journal of the Neurological Sciences* 269: 96–104.

McMahon, D. B., Leopold, D. A. (2012). Stimulus timing-dependent plasticity in highlevel vision. *Current Biology* 22: 332–337.

McNickle, E., Carson, R. G. (2015). Paired associative transcranial alternating current stimulation increases the excitability of corticospinal projections in humans. *Journal of Physiology* 593: 1649-66.

Mele, A. (2014). *Free: Why science hasn't disproved free will*. Oxford: Oxford University Press.

Mrachacz-Kersting, N., Kristensen, S. R., Niazi, I. K., Farina, D. (2012). Precise temporal association between cortical potentials evoked by motor imagination and afference induces cortical plasticity. *Journal of Physiology* 590: 1669-82.

Nazzaro, C., Greco, B., Cerovic, M., Baxter, P., Rubino, T., Trusel, M., Parolaro, D., Tkatch, T., Benefenati, F., Pedarzani, P., Tonini, R. (2012). SK channel modulation rescues striatal plasticity and control over habit in cannabinoid tolerance. *Nature Neuroscience* 15: 284-93.

O'Connor, T. (2000). *Persons and causes: The metaphysics of free will*. Oxford: Oxford University Press.

O'Connor, T. (2008). 'Agent-causal power', in *Dispositions and causes*, ed. T. Handfield. Oxford: Oxford University Press.

Orban de Xivry, J. J., Ahmadi-Pajouh, M. A., Harran, M. D., Salimpour, Y., Salimpour, Y., Shadmehr, R. (2013). Changes in corticospinal excitability during reach adaptation in force fields. *Journal of Neurophysiology* 109: 124-36.

Palva JM, Zhigalov A, Hirvonen J, Korhonen O, Linkenkaer-Hansen K, Palva S. (2013).

Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proceedings from the National Academy of Science USA* 110: 3585-90.

Pereboom, D. (1995). Determinism al dente. *Noûs* 29: 21-45.

Pereboom, D. (2001). *Living without free will*. Cambridge: Cambridge University Press.

Pereboom, D. (2013). 'Free will skepticism and criminal punishment', in *The future of punishment*, ed. T. Nadelhoffer. New York: Oxford University Press.

Roy, F. D., Norton, J. A., Gorassini, M. A. (2007). Role of sustained excitability of the leg motor cortex after transcranial magnetic stimulation in associative plasticity. *Journal of Neurophysiology* 98: 657-67.

Dash, P. K., Hebert, A. E., Runyan, J. D. (2004). A unified theory for systems and cellular memory consolidation. *Brain Research Reviews* 45: 30-7.

Runyan, J. D., (2016). Events, agents, and settling whether and how one intervenes. *Philosophical Studies* 173: 1629. doi:10.1007/s11098-015-0572-4.

Shannon, B. J., Dosenbach, R. A., Su, Y., Vlassenko, A. G., Larson-Prior, L., J., Nolan, T. S., Snyder, A. Z., Raichle, M. E. (2013). Morning-evening variation in human brain metabolism and memory circuits. *Journal of Neurophysiology* 109: 1444-56.

Schurger, A., Sitt, J. D., and Dehaene, S. (2012). An accumulator model for spontaneousneural activity prior to self-initiated movement. *Proceedings from the National Academy of Science U S A* 109: E2904–E2913.

Schurger, A., and Uithol, S. (2015). Nowhere and everywhere: the causal origin of voluntary action. *The Review of Philosophy and Psychology* 6: 761–778.

Schurger, A., Mylopoulos, M., and Rosenthal, D. (2016). Neural antecedents of spontaneous voluntary movement: a new perspective. *Trends in Cognitive Science* 20: 77–79.

Smith, K. S., Graybiel, A. M., (2013). A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron* 79: 361-74.

Steward, H. (2012). *A metaphysics for freedom*. Oxford: Oxford University Press.

Svensson E., Evans C. G., Cropper E. C. (2016). Repetition priming-induced changes in sensorimotor transmission. *Journal of Neurophysiology* 115: 1637-43.

Vargas, M. (2013). *Building better beings: A theory of moral responsibility* (Oxford: Oxford University Press).

Veniero, D., Ponzo, V., Koch, G. (2013). Paired associative stimulation enforces the communication between interconnected areas. *Journal of Neuroscience* 33: 13775-83.

Verstynen, T., Sabes, P. N. (2011). How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *Journal of Neuroscience* 31: 10050-9.

Wang, X.-J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews* 90: 1195-1268.

Wickens, J. R., Horvitz, J. C., Costa, R. M., Killcross, S., (2007). Dopaminergic mechanisms in actions and habits. *Journal of Neuroscience* 27: 818-3.

Wolters, A., Sandbrink, F., Schlottmann, A., Kunesch, E., Stefan, K., Cohen, L. G., Benecke, R., Classen, J. (2003). A temporally asymmetric Hebbian rule governing plasticity in the human motor cortex. *Journal of Neurophysiology* 89: 2339–2345.

Yin, H. H., Mulcare, S. P., Hilário, M. R., Clouse, E., Holloway, T., Davis, M. I., Hansson, A. C., Lovinger, D. M., Costa, R. M. (2009). Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nature Neuroscience* 12: 333-41.